Voting for Deceptive Opinion Spam Detection

Tao Wang and Hua Zhu
Computer Science Department, Nanjing University
Wuhan, Hubei, P.R. China, 430072
pwangtao@gmail.com, huaz@whu.edu.cn

arXiv:1409.4504v1 [cs.CL] 16 Sep 2014

Abstract

Consumers' purchase decisions are increasingly influenced by user-generated online reviews. Accordingly, there has been growing concern about the potential for posting deceptive opinion spam fictitious reviews that have been deliberately written to sound authentic, to deceive the readers. Existing approaches mainly focus on developing automatic supervised learning based methods to help users identify deceptive opinion spams. this work, we used the LSI and Sprinkled LSI technique to reduce the dimension for deception detection. We make our contribution to demonstrate what LSI is capturing in latent semantic space and reveal how deceptive opinions can be recognized automatically from truthful opinions. Finally, we proposed a voting scheme which integrates different approaches to further improve the classification performance.

1 Introduction

Consumers increasingly rely on user-generated online reviews when making purchase decision (Cone, 2011; Ipsos, 2012). Unfortunately, the ease of posting content to the Web, potentially anonymously, create opportunities and incentives for unscrupulous businesses to post deceptive opinion spam fraudulent or fictitious reviews that are deliberately written to sound authentic, in order to deceive the reader (Ott et al., 2011). For example, a hotel manager may post fake positive reviews to promote their own hotel, or fake negative reviews to demote a competitor's hotel. Accordingly, there appears to be widespread and growing concern among both businesses and the public (Meyer, 2009; Miller, 2009; Streitfeld, 2012; Topping, 2010). The followings are two reviews, one is deceptive and the other one is truthful.

1. My husband and I arrived for a 3 night stay. for our 10th wedding anniversary. We had booked an Executive Guest room, upon ar- rival we were informed that they would be upgrading us to a beautiful Junior Suite. This was just a wonderful unexpected plus to our beautifully planned weekend. The front desk manager was professional and made us feel warmly welcomed. The Chicago Affinia was just a gorgeous hotel, friendly staff, lovely food and great atmosphere. Not the mention the feather pillows and bedding that was just fantastic. Also we were allowed to bring out beloved Shi-Tzu and he experienced the Jet Set Pets stay. The grooming was perfect, the daycare service we felt completely comfortable with. This was a beautiful weekend, thank you Affinia Hotels! We would visit this hotel again!

2. *As others have said, once all the construction works are completed I suspect that the prices this hotel can (legitimately) charge will put it out of our price range. Which will be a pity because it is an excellent hotel, the location couldn't be better. The room was very spacious, with separate sitting study areas a nice bathroom. Only 3 minor points: downstairs they were serving a free basic breakfast (coffee and pastries), but we only knew of it on our last morning nobody had mentioned this; the cost of internet is a bit dear especially when lots of motels now offer it free; there was only powered milk in the rooms, which wasn't that nice. But none of these really spolied a really enjoyable stay.*.

As can be seen, as the distinction is too subtle, it is hard to manually tell truthful reviews from deceptive one (the first one is deceptive). Existing approaches for spam detection usually focus on developing supervised learning-based algo-

rithms to help users identify deceptive opinion spam (Jindal and Liu, 2008; Jindal et al., 2010; Li et al., 2011; Lim et al., 2010; Ott et al., 2011; Wang et al., 2011; Wu et al., 2010). These supervised approaches suffer one main disadvantage: they are highly dependent upon high-quality gold-standard labelled data. One option for producing gold-standard labels, for example, would be to rely on the judgements of human annotators (Jindal et al., 2010; Mukherjee et al., 2012). Recent studies, however, shows that unlike other kinds of spam, such as Web (Castillo et al., 2006; Martinez-Romo and Araujo, 2009) and e-mail spam (Chirita et al., 2005), deceptive opinion spam is neither easily ignored nor easily identified by human readers (Ott et al., 2011). This is especially the case when considering the overtrusting nature of most human judges, a phenomenon referred to in the psychological deception literature as a truth bias (Vrij, 2008).

Due to the difficulty in manually labeling deceptive opinion, Ott et al. (2011) (Ott et al., 2011) solicit deceptive reviews from workers on Amazon Mechanical Turk, and built a dataset containing 400 deceptive and 400 truthful reviews, which they use to train and evaluate supervised SVM classifiers[1]. According to their findings, truthful hotel opinions are more specific about spatial configurations(e.g. small, bathroom, location), which can be easily explained by the fact that the Turkers have never been to that hotels.

Our work in this paper started from the dataset, 400 truthful and 400 deceptive reviews, from (Ott et al., 2011; Li et al., 2013a), a large-scale, publicly available dataset for deceptive opinion spam research. Then we use three ways to preprocessing dataset. (a) Use N-gram language model to build a probabilistic language model (b) Use POS-tag to delve the linguistic reason why a review would likely to be considered fake (c) Use TF-IDF to reflect how important a word is in given text. After pre-processing, the data has high dimensional features which is quite computational expensive and easily raises the over fitting problem. Thus we are motivated to reduce dimension using Latent Semantic Indexing (LSI), one of most popular indexing and retrieval method that uses singular value decomposition (SVD) to map the documents into vector in the latent concept-space. LSI extracts useful information in latent concepts through learning a low-dimensional representation of this text, which helps us to know more detailed difference in fake and real reviews. We also use supervised latent semantic analysis (Sprinkle) to overcome the shortcoming of LSI and results show performance improves a lot. To classify data, we use Support Vector Machine (SVM) and Naive Bayes, two popular and mature classifier, on our dataset after pre-processing and dimension reduction. Although these classifiers are quite powerful, it may still misclassify correct points. So we use a voting scheme or a weighted combination of the multiple hypothesis to achieve final conclusion. The remainder of this paper is organized as follows: Section 2 talks about related work. Section 3 presents details of our fake reviews detection framework, including three approaches to preprocessing data, two ways of dimension reduction, and two classification method. In section 4 shows the experimental results and discussion of performance; Finally, Section 5 concludes our work.

[1]Truthful opinions are selected from 5-star reviews from TripAdvisor. http://www.tripadvisor.com/

2 Related Work

Jindal and Liu (2008) first studied the deceptive opinion problem and trained models using features based on the review text, reviewer, and product to identify duplicate opinions, i.e., opinions that appear more than once in the corpus with similar contexts. Wu et al. (2010) propose an alternative strategy to detect deceptive opinion spam in the absence of a gold standard. Yoo and Gretzel (2009) gathered 40 truthful and 42 deceptive hotel reviews and manually compare the linguistic differences between them. Ott et al. created a gold-standard collection by employing Turkers to write fake reviews, and follow-up research was based on their data (Ott et al., 2012; Ott et al., 2013; Li et al., 2013b; Feng and Hirst, 2013; Li et al., 2014). For example, Song et al. (2012) looked into syntactic features from Context Free Grammar parse trees to improve the classifier performance. A step further, Feng and Hirst (2013) make use of degree of *compatibility* between the personal experiment and a collection of reference reviews about the same product rather than simple textual features.

In addition to exploring text or linguistic features in deception, some existing work looks into customers' behavior to identify deception (Mukherjee et al., 2013). For example, Mukherjee

et al. (2011; 2012) delved into group behavior to identify group of reviewers who work collaboratively to write fake reviews.

3 Data Processing

Before building n-gram model, we tokenized the text of each review and used porter stemming algorithms to eliminate the influence of different forms of one word. An n-gram model is a type of 2 probabilistic language model for predicting the next item in such a sequence in the form of n-1 order Markov model. We built unigram and bigram model in our experiment.

tf-idf TF-IDF, short for term frequencyinverse document frequency, is a numerical statistic that is intended to reflect how important a word is to a document in a collection or corpus. The tf-idf value increases proportionally to the number of times a word appears in the document, but is offset by the frequency of the word in the corpus, which helps to control for the fact that some words are generally more common than others. For example, stop words – common words that appears every where but do not have much meaning, such as the, a, and that, are of little significance. So tf-idf will assign small weight to these words.

The formulae of tf-idf is:

$$w = tf \cdot idf_{i,j} = tf \cdot \log \frac{N}{df_{i,j}}$$

POS Part-of-speech tagging, also called grammatical tagging or word-category disambiguation, is the process of marking up a word in a text as corresponding to a particular part of speech. Our project use POS-tag to delve the linguistic reason why a review would likely to be considered fake. We use the Stanford Parser to obtain the relative POS frequencies as feature vector to separate data.

4 Model

The power of dimension reduction lies in the fact that it can ease later processing, improve computational performance improvement, filter useless noise and recover underlying causes. In our project, we are using LSI for text classification. LSI is based upon the assumption that there is an underlying semantic structure in textual data, and that the relationship between terms and documents can be re-described in this semantic structure form. Textual documents are represented as vectors in a vector space. It is fundamentally based on SVD that breaks original relationship of the data into linearly independent components, where the original term vectors are represented by left singular vectors and document vectors by right singular vectors. That is

$$U_{d\times l} s_{l\times l} V_{l\times d} \tag{1}$$

The column of $U_{d\times l}$ defines the lower dimensional coordinate system. S_l is a diagonal matrix with the l largest singular values in non-increasing order along its diagonal. The l columns are the new coordinates of each document after dimensionality reduction.

One major key of LSI is that it overcomes two of the most problematic constraints synonymy and polysemy. It is very tolerant of noise (i.e, misspelled words, typographical errors, etc.) and has been proven to capture key relationship information, including causal, goal-oriented, and taxonomic information.Another usage of LSI is in data visualization. Let l=2 or 3. Each xi is approximated by a l-dimensional vector now suitable for plotting in 2D or 3D spaces.

LSI has limitations in classification because it is an unsupervised method which doesn't take class information into account. So we decide to use the process Sprinkling in which LSI is performed on a term-document matrix augmented with sprinkled terms, namely class labels.

4.1 Classification

4.2 Naive Bayes

Naive Bayes classifiers are among the most successful known algorithms for learning to classify text documents. As we know, text classifiers often don't use kind of deep representation about language: often a document is represented as a bag of words. This is a very simple representation of document : it only knows which words are included in the document and throws away the word order. NB classifier just relies on this representation. For a document D and a class C. In our project, we only have two classes: fake or real. We classify d as the class which has the highest posterior probability P(C—D), which can be re-expressed using Bayes Theorem:

$$p(c|d) = \frac{p(d|c)p(c)}{p(d)} \tag{2}$$

$$
\begin{aligned}
C_{MAP} &= argmax_c P(c|d) \\
&= argmax_c \frac{P(d|c)P(c)}{p(d)} \\
&= argmax_c P(d|c)P(c)
\end{aligned} \tag{3}
$$

4.3 Support Vector Machine

Support Vector Machines(SVM) (Joachims, 1999) are supervised learning algorithms that used to recognize patterns and classify data. Given a set of binary training data, an SVM training algorithm builds a model to calculate a hyper-plane that separate data into one category or the other. The goal of SVM is to find a hyper-plane that clearly classify data and maximize the distance of the support vectors to hyper-plane at the same time .

$$
\frac{1}{2}||w||^2 - \sum_i \alpha_i [y_i(w \cdot x_i - b) - 1] \tag{4}
$$

We choose SVM as our classifier since in the area of text classification, SVM performs more robust than other techniques. When implementing SVM, we first pre-process each original document to a vector of features. Such text feature vectors often have high dimension, however, SVM supports quite well. SVMs use overfitting protection, so the performance does not affect much when the number of features increases. Joachims shows most text categorization problems are linearly separable, so we use linear SVMs, which is also faster to learn, to solve our problem. In our project, we use SVMlight to train our linear SVM models on datasets that processed by UNIGRAMS, BIGRAMS, POS-Tag, and TF-IDF. We also train dataset after dimension reduction. Results show that SVMs perform well on classifying our data.

4.4 Voting

Voting Scheme, also called weighted combination of the multiple hypothesis, is used to improve performance when having different approaches to a certain problem. Although some classifier is quite powerful, it may still misclassify correct points. On the other hand, we desire simple models that free users from troublesome algorithm design. The uncertainty of dimension selection and goodness evaluation makes classification difficult. This method is more robust to difficult classified text and achieves better performance comparing to separate models.

5 Experimental Results

To better understand the models learned by these automated approaches, we decide to demonstrate what SVD in LSI analysis is capturing. We find the group of words that have the most significant influence in latent concepts. we can see, the second concept contributes most for separating the fake reviews. We list in Table with the highest scores that loads on the positive polar (truthful) and negative polar (fake) of second concept.

deceptive	truthful
hotel	room
my	)
chicago	(
will	but
room	$
very	bathroom
visit	location
husband	night
city	walk
experience	park

Table 1: Top words in different topics from Topic-Spam

We are using Sprinkled LSI and LSI to reduce the term-document matrix in different dimensions from 50 to 700. We use SVM to classify the data. The results show that with the dimension increasing, the training accuracy of two methods is also increasing because we have higher features which give us more information. Besides, the testing accuracy and precision is increasing when dimension is from 0 to 500. When dimension continues to increase, the training accuracy increases while testing accuracy goes down. It implies to us that there is an overfitting problem when the dimension of the training data is high. Results in bold responds to best accuracy. We can see that when we use Sprinkle + SVM and the dimension is around 500, the accuracy is 90% which is quite good.

In this paper, we combine 5 different algorithms above to automatically obtain the final classification by voting for the test results. If over 3 methods vote for the same class, we will determine the document to belong to that class. The five algorithms we choose as follows: Sprinkle+SVM (dimension 500), Sprinkle+SVM(dimension 300), Unigram+SVM, TF.IDF+SVM, Unigram+NB. The voting algorithm gets 0.95 accuracy.

$\text{SVM}_{unigram}$	0.90
$\text{NB}_{unigram}$	0.863
SVM_{tf-idf}	0.885
$\text{SVM}_{sprinkle300}$	0.890
$\text{SVM}_{sprinkle500}$	0.900
Voting	0.95

Table 2: Performance comparisons of Voting Scheme vs separate algorithm

6 Conclusion

In this work we use different methods to analyse and improve fake reviews classification result based on Myle Ott and Yejin Choi work. It shows that the detection of fake reviews by using automatically approaches is well beyond using human judges. We first build n-gram models and POS model combined with SVM and Naive Bayes to detect the deceptive opinion. Similar to Myle Ott and Yejin Chois result, n-gram based test categorization have the best performance suggesting that key-word based approaches might be necessary in detecting fake review. Besides, we used LSI to exert dimension reduction and some theoretical analysis. Contrast to Myle Ott and Yejin Choi, our result shows that 2nd person pronouns, rather than 1st person form are relatively tend to be used in deceptive review. The fake reviews also have the character of lacking of the concrete noun and adjectives. Furthermore, we compare the experimental results of Sprinkled LSI and LSI . We find Sprinkled LSI+ SVM outperforms LSI + SVM when dimension is not high. Finally, we proposed a voting scheme which combines a few algorithms we used above to achieve better classification performance.

References

[Castillo et al.2006] Carlos Castillo, Debora Donato, Luca Becchetti, Paolo Boldi, Stefano Leonardi, Massimo Santini, and Sebastiano Vigna. 2006. A reference collection for web spam. In *ACM Sigir Forum*, volume 40, pages 11–24. ACM.

[Chirita et al.2005] Paul-Alexandru Chirita, Jörg Diederich, and Wolfgang Nejdl. 2005. Mailrank: using ranking for spam detection. In *Proceedings of the 14th ACM international conference on Information and knowledge management*, pages 373–380. ACM.

[Cone2011] Cone. 2011. 2011 Online Influence Trend Tracker. http://www.coneinc.com/negative-reviews-online-reverse-purchase-decisions, August.

[Feng and Hirst2013] Vanessa Wei Feng and Graeme Hirst. 2013. Detecting deceptive opinions with profile compatibility. In *Proceedings of the 6th International Joint Conference on Natural Language Processing, Nagoya, Japan*, pages 14–18.

[Feng et al.2012] Song Feng, Ritwik Banerjee, and Yejin Choi. 2012. Syntactic stylometry for deception detection. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers-Volume 2*, pages 171–175. Association for Computational Linguistics.

[Ipsos2012] Ipsos. 2012. Socialogue: Five Stars? Thumbs Up? A+ or Just Average? http://www.ipsos-na.com/news-polls/pressrelease.aspx?id=5929.

[Jindal and Liu2008] Nitin Jindal and Bing Liu. 2008. Opinion spam and analysis. In *Proceedings of the international conference on Web search and web data mining*, pages 219–230. ACM.

[Jindal et al.2010] Nitin Jindal, Bing Liu, and Ee-Peng Lim. 2010. Finding unusual review patterns using unexpected rules. In *Proceedings of the 19th ACM international conference on Information and knowledge management*, pages 1549–1552. ACM.

[Joachims1999] Thorsten Joachims. 1999. Making large scale svm learning practical.

[Li et al.2011] Fangtao Li, Minlie Huang, Yi Yang, and Xiaoyan Zhu. 2011. Learning to identify review spam. In *Proceedings of the Twenty-Second international joint conference on Artificial Intelligence-Volume Volume Three*, pages 2488–2493. AAAI Press.

[Li et al.2013a] Jiwei Li, Claire Cardie, and Sujian Li. 2013a. Topicspam: a topic-model based approach for spam detection. In *ACL (2)*, pages 217–221.

[Li et al.2013b] Jiwei Li, Myle Ott, and Claire Cardie. 2013b. Identifying manipulated offerings on review portals. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, Seattle, Wash*, pages 18–21.

[Li et al.2014] Jiwei Li, Myle Ott, Claire Cardie, and Eduard Hovy. 2014. Towards a general rule for identifying deceptive opinion spam.

[Lim et al.2010] Ee-Peng Lim, Viet-An Nguyen, Nitin Jindal, Bing Liu, and Hady Wirawan Lauw. 2010. Detecting product review spammers using rating behaviors. In *Proceedings of the 19th ACM international conference on Information and knowledge management*, pages 939–948. ACM.

[Martinez-Romo and Araujo2009] Juan Martinez-Romo and Lourdes Araujo. 2009. Web spam identification through language model analysis.

In *Proceedings of the 5th International Workshop on Adversarial Information Retrieval on the Web*, pages 21–28. ACM.

[Meyer2009] David Meyer. 2009. Fake reviews prompt belkin apology.

[Miller2009] C Miller. 2009. Company settles case of reviews it faked. *New York Times*.

[Mukherjee et al.2011] Arjun Mukherjee, Bing Liu, Junhui Wang, Natalie Glance, and Nitin Jindal. 2011. Detecting group review spam. In *Proceedings of the 20th international conference companion on World wide web*, pages 93–94. ACM.

[Mukherjee et al.2012] Arjun Mukherjee, Bing Liu, and Natalie Glance. 2012. Spotting fake reviewer groups in consumer reviews. In *Proceedings of the 21st international conference on World Wide Web*, pages 191–200. ACM.

[Mukherjee et al.2013] Arjun Mukherjee, Abhinav Kumar, Bing Liu, Junhui Wang, Meichun Hsu, Malu Castellanos, and Riddhiman Ghosh. 2013. Spotting opinion spammers using behavioral footprints. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 632–640. ACM.

[Ott et al.2011] Myle Ott, Yejin Choi, Claire Cardie, and Jeffrey T. Hancock. 2011. Finding deceptive opinion spam by any stretch of the imagination. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 309–319.

[Ott et al.2012] Myle Ott, Claire Cardie, and Jeff Hancock. 2012. Estimating the prevalence of deception in online review communities. In *Proceedings of the 21st international conference on World Wide Web*, pages 201–210. ACM.

[Ott et al.2013] Myle Ott, Claire Cardie, and Jeffrey T. Hancock. 2013. Negative deceptive opinion spam. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Short Papers*, Atlanta, Georgia, USA, June. Association for Computational Linguistics.

[Streitfeld2012] David Streitfeld. 2012. For 2 a star, an online retailer gets 5-star product reviews. *New York Times.*, 26.

[Topping2010] A Topping. 2010. Historian orlando figes agrees to pay damages for fake reviews. *The Guardian.*, 16.

[Vrij2008] Aldert Vrij. 2008. *Detecting lies and deceit: Pitfalls and opportunities*. Wiley. com.

[Wang et al.2011] Guan Wang, Sihong Xie, Bing Liu, and Philip S Yu. 2011. Review graph based online store review spammer detection. In *Data Mining (ICDM), 2011 IEEE 11th International Conference on*, pages 1242–1247. IEEE.

[Wu et al.2010] Guangyu Wu, Derek Greene, Barry Smyth, and Pádraig Cunningham. 2010. Distortion as a validation criterion in the identification of suspicious reviews. In *Proceedings of the First Workshop on Social Media Analytics*, pages 10–13. ACM.

[Yoo and Gretzel2009] Kyung-Hyan Yoo and Ulrike Gretzel. 2009. Comparison of deceptive and truthful travel reviews. In *Information and communication technologies in tourism 2009*, pages 37–47. Springer.